arXiv:1604.07342v2 [cs.CV] 9 Jun 2016

Supervised Incremental Hashing

Bahadir Ozdemir[1]
ozdemir@cs.umd.edu
Mahyar Najibi[1]
najibi@cs.umd.edu
Larry S. Davis[2]
lsd@umiacs.umd.edu

[1] Department of Computer Science
University of Maryland
College Park, MD 20742

[2] Institute for Advanced Computer Studies
University of Maryland
College Park, MD 20742

Abstract

We propose an incremental strategy for learning hash functions with kernels for large-scale image search. Our method is based on a two-stage classification framework that treats binary codes as intermediate variables between the feature space and the semantic space. In the first stage of classification, binary codes are considered as class labels by a set of binary SVMs; each corresponds to one bit. In the second stage, binary codes become the input space of a multi-class SVM. Hash functions are learned by an efficient algorithm where the NP-hard problem of finding optimal binary codes is solved via cyclic coordinate descent and SVMs are trained in a parallelized incremental manner. For modifications like adding images from a previously unseen class, we describe an incremental procedure for effective and efficient updates to the previous hash functions. Experiments on three large-scale image datasets demonstrate the effectiveness of the proposed hashing method, Supervised Incremental Hashing (SIH), over the state-of-the-art supervised hashing methods.

1 Introduction

Online image databases like Flickr grow steadily every day. Huge amounts of information lead to the requirement of efficient search algorithms. The linear time complexity of exhaustive similarity search can be reduced to sublinear time complexity for approximate nearest neighbor search through discretization [5]. Several approaches have been proposed to discretize the feature space by preserving pairwise similarities between data points. Binary hashing methods, which map data points to short binary codes, have attracted attention from researchers in different areas due to their effectiveness in large-scale search and efficiency in memory storage. The methods for learning hash functions are categorized into unsupervised and supervised learning schemes. In this paper, we only focus on supervised hashing.

Supervised hashing methods are constructed to capture semantic similarity between images. Semantic information such as class labels or tags is used to regularize the process of learning binary codes. Several supervised methods have been proposed as extensions to unsupervised techniques by making them aware of the semantics of data. Iterative Quantization (ITQ) is used for both unsupervised and supervised hashing in conjunction with Principal Component Analysis (PCA) and Canonical Correlation Analysis (CCA), respectively [6]. A

© 2016. The copyright of this document resides with its authors.
It may be distributed unchanged freely in print or electronic forms.

supervised version of BRE was proposed by Kulis and Darrell [10]. Supervised Hashing with Kernels (KSH) [15] was introduced based on the same approach towards hashing. Discriminative binary coding methods [8, 17], like Random Maximum Margin Hashing approach, try to improve generalization by using SVMs to preserve the semantic structure in the Hamming space. Lin *et al.* [13] proposed a supervised hashing method called FastHash that tries to capture the non-linearity in the feature space utilizing decision trees and uses a GraphCut based method for optimization. Supervised discrete hashing (SDH) [18] employs a linear support vector machine (SVM) to capture the similarity in the semantic space.

Despite the fact that new images are added to online photo databases every day, to the best of our knowledge, no supervised hashing method learns hash functions incrementally for newly added images. Efficiently updating hash functions with new images is a challenging task. Besides, some of the new images will possibly not belong to existing classes *i.e.* forming new semantic classes, that makes this task even more challenging. For the existing hashing methods, such modifications to a database require recomputation of hash functions from scratch whenever a change occurs in a dynamic dataset, which is computationally intractable.

In this paper, we propose an incremental supervised hashing method with kernels based on binary and multi-class SVMs, which we refer to as *Supervised Incremental Hashing* (SIH). We adopt the incremental learning fashion of SVMs in a hashing framework. Consequently, SIH can be easily extended to unseen classes incrementally. We identified three main objectives for our supervised hashing method – as being incremental and parallelizable, avoiding overfitting by better generalization, and balancing $+1/-1$ in learned binary codes. The significance of balanced binary codes has been studied in the hashing literature [7, 11, 14, 20] and their absence leads to ineffective codes, especially when the code length is small and/or the sizes of semantic classes are unbalanced. To overcome these issues, we reformulate the supervised hashing task as a two-stage classification framework. In the first stage of classification, we use a binary SVM for each bit. The feature space is the input space of SVMs while binary codes are considered as class labels. In the second stage, the binary codes become the input space of a single multi-class SVM, and the semantic space is the output space of the SVM. Also, we penalize imbalance in binary codes in our optimization formulation.

We define an incremental strategy to learn SVMs and a *discrete cyclic coordinate descent* (DCC) algorithm, similar to the one of SDH, to learn binary codes bit by bit as an approximate solution to the NP-hard discrete optimization problem. Our contributions can be summarized as follows:

- We propose a supervised hashing approach that provides better generalization with regularizations and maximizes the entropy by balancing binary codes. We formulate our hashing objectives in a single optimization task.
- We describe an algorithm that solves the optimization problem efficiently by an incremental strategy for training SVMs and an approximation to the solution of an NP-hard problem.
- We define an incremental strategy for the proposed hashing method that takes the earlier hash functions and the final state of the database as its input and efficiently computes new hash functions which perform similarly to those computed from scratch on the final state of the database.

The rest of the paper is organized as follows: Section 2 introduces the proposed hashing method. Section 3 gives performance evaluation and comparison with the state-of-the-art supervised hashing methods, and Section 4 provides conclusions.

2 Incremental Hashing

2.1 Problem Definition

Given a set of data points $\{\mathbf{x}_i\}_{i=1}^n$, each $\mathbf{x}_i \in \mathbb{R}^d$, the goal of supervised hashing is to learn a function H that maps data points to m-dimensional binary codes capturing underlying semantics. Several hashing methods [5, 6] construct such a hash function by collecting a set of m linear binary embeddings $H(\mathbf{x}) = [h_1(\mathbf{x}), \dots, h_m(\mathbf{x})]^\top$ such that $h_j(\mathbf{x}) = \text{sign}(\mathbf{w}_j^\top \mathbf{x} + \beta_j)$ where $\mathbf{w}_j \in \mathbb{R}^d$ is the normal vector and $\beta_j \in \mathbb{R}$ is the intercept of a hyperplane for $j = 1, \dots, m$. Some methods [11, 15] employ a nonlinear embedding algorithm with a kernel function $\varphi : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ such that each binary embedding has the form $h_j(\mathbf{x}) = \text{sign}(\mathbf{w}_j^\top \varphi(\mathbf{x}) + \beta_j)$ where $\varphi(\mathbf{x}) = [\varphi(\mathbf{x}, \mathbf{a}_1), \dots, \varphi(\mathbf{x}, \mathbf{a}_r)]^\top$ and $\{\mathbf{a}_l\}_{l=1}^s$ is a set of anchor points randomly selected from training samples.

2.2 Learning Hash Functions for Dynamic Databases

Modifications to an image database consist of adding new images to the database and deleting images from the database. Considering a database with class information, there are four types of modifications:

1. Adding new classes,
2. Adding images to existing classes,
3. Deleting existing classes,
4. Deleting images from existing classes.

In this paper, we focus on the first three types of modifications. From a hashing perspective, the last type of modifications is the least interesting one because it technically results in reducing training size. Among the rest, deleting existing classes is the easiest case because we have already learned binary codes for all images in the final database. On the other hand, adding new classes is the hardest case since we have no prior information about the new classes.

We define three hashing strategies for dynamic datasets:

1. *Passive strategy*: Continuing to use the hash functions learned from an earlier state of the database *i.e.* ignoring the changes in the training data.
2. *From-scratch strategy*: Learning hash functions on the final state of the database from scratch whenever a change occurs in the database.
3. *Incremental strategy*: Learning hash functions on the last state of the database incrementally starting from the previous hash functions.

The incremental strategy can be considered as effective and efficient if its retrieval performance is similar to that of the from-scratch strategy while its training time is shorter. We first describe our supervised hashing method; then provide an incremental hashing strategy for our method in the following sections.

2.3 Supervised Incremental Hashing

We define a joint optimization task as a combination of m binary linear SVMs and a multi-class linear SVM by Crammer and Singer [3] and an imbalance penalty. Our method searches

for optimal set of binary codes $\{\mathbf{b}_i\}_{i=1}^n$ where $\mathbf{b}_i = [b_{i1}, \ldots, b_{im}]^\top$ corresponds to the ith data point in the training set. Let $\mathcal{Y}$ denote the set of ground truth semantic classes and $y_i \in \mathcal{Y}$ denote the class label of the ith data point; then we define a multi-class SVM model that maps binary codes to semantic classes *i.e.* $\mathbf{b}_i \mapsto y_i$ for $i = 1, \ldots, n$.

Let $\mathbf{B} \in \{-1, +1\}^{n \times m}$ be a matrix that collects binary codes where the ith row of $\mathbf{B}$ is equal to $\mathbf{b}_i^\top$. We define a binary SVM model treating $\{\varphi(\mathbf{x}_i)\}_{i=1}^n$ as input data and column j of $\mathbf{B}$ as binary class labels for $j = 1, \ldots, m$. In addition, we add a new term that penalizes imbalanced assignments of $+1/-1$ in binary codes. Our formulation for supervised hashing is given by

$$\begin{aligned} \min_{\mathbf{B}, \mathcal{W}, \Xi} \quad & \lambda \left(\frac{1}{2} \sum_{k \in \mathcal{Y}} \|\mathbf{w}_k^b\|^2 + C^b \sum_{i=1}^n \xi_i^b \right) + \sum_{j=1}^m \left(\frac{1}{2} \|\mathbf{w}_j^x\|^2 + C^x \sum_{i=1}^n \xi_{ij}^x \right) + \gamma \sum_{j=1}^m \left| \sum_{i=1}^n b_{ij} \right| \\ \text{s.t.} \quad & \forall (i,k) \quad (\mathbf{w}_{y_i}^b - \mathbf{w}_k^b)^\top \mathbf{b}_i \geq \mathbb{1}[y_i \neq k] - \xi_i^b, \\ & \forall (i,j) \quad b_{ij} (\mathbf{w}_j^x)^\top \varphi(\mathbf{x}_i) \geq 1 - \xi_{ij}^x, \\ & \forall (i,j) \quad \xi_{ij}^x \geq 0, \\ & \mathbf{b}_i \in \{-1, +1\}^m \end{aligned} \tag{1}$$

where λ and γ are scaling parameters, C^x and C^b are soft margin parameters of the SVMs. Note that the sum of column j in $\mathbf{B}$ is equal to zero when bit j is balanced. There will be no penalty for bit j in this case. This imbalance penalty is also important for the performance of the binary SVMs with the fixed parameter C^x since they are sensitive to imbalance in the datasets [1]. We can add an extra element 1 to the vectors of all data points and binary codes in order to add a bias term to the loss function of SVMs. Finally, our hash function is defined as $H(\mathbf{x}) = \text{sign}\big((\mathbf{W}^x)^\top \varphi(\mathbf{x})\big)$ where the jth column of $\mathbf{W}^x \in \mathbb{R}^{d \times m}$ is equal to $\mathbf{w}_j^x$ for $j = 1, \ldots, m$. Similarly, we define another weight matrix $\mathbf{W}^b \in \mathbb{R}^{d \times m}$ such that its jth column is equal to $\mathbf{w}_k^b$ for $k = 1, \ldots, |\mathcal{Y}|$. This optimization task can be solved by an alternating algorithm that consists of two main steps: Training SVMs and learning binary codes as follows.

2.4 Training SVMs

If we fix all binary codes $\mathbf{B}$, the optimization problem (1) will be reduced to m separate binary SVM and one multi-class linear SVM problems. These $(m+1)$ SVMs can be trained in parallel since they are conditionally independent on binary codes. We employ an efficient large-scale SVM learning technique called the *Optimized Cutting Plane Algorithm for SVMs* (OCAS) [2]. The OCAS algorithm is based on approximating the original convex loss function of an SVM by a piecewise linear function defined as the maximum over a set of linear under-estimators called cutting planes. At each step, a new cutting plane is added to the set with a cost of $\mathcal{O}(nd)$ complexity. Each under-estimator is a linear approximation to the loss function of the SVM at a weight point $\mathbf{w}'$. The method also has a multi-class counterpart called OCAM.

The OCAS method tries to minimize the total number of iterations by additional computations at each step. Let $F(\mathbf{w}) = \frac{1}{2}\|\mathbf{w}\|^2 + CR(\mathbf{w})$ be an SVM objective function such that $\|\partial R(\mathbf{w})\| \leq G$, then the OCAS method converges after at most

$$\log_2 \frac{F(\mathbf{0})}{4C^2G^2} + \frac{8C^2G^2}{\varepsilon} - 2$$

iterations where ε denotes a tolerance value. We can further reduce the total number iterations for an SVM with hinge loss by $\log_2 \frac{nC}{F(\mathbf{w}_0)}$ if there exists an initial weight vector $\mathbf{w}_0$ better than the zero vector *i.e.* $F(\mathbf{w}_0) < F(\mathbf{0})$ (See Appendix). Considering the cases of adding new images to existing classes and deleting classes, the solution to the previous SVM problem can be a good starting point. In addition, changes in binary codes $\mathbf{B}$ from one iteration to the next one are usually small during the execution of our algorithm. Therefore, we take advantage of the pre-computed SVMs from the previous iteration in an incremental approach. We employ a warm start strategy for training SVMs similar to [19]. We initialize the *best-so-far solution* in the OCAS algorithm with the solution to the SVM in the previous iteration instead of a zero vector for each bit. This initialization results in a reduced number of iterations in the OCAS computation and fewer cutting planes. If there exists no change in column j of $\mathbf{B}$ between iteration $(t-1)$ and iteration t, the corresponding SVM is not trained again. Similarly, the same approach is used for the multi-class linear SVM using the OCAM algorithm. Our hashing method converges when the entire matrix $\mathbf{B}$ remains unchanged. A simulation of our incremental SVM approach is presented on a sample dataset in Figure 1.

2.5 Learning Binary Codes

In the process of learning binary codes $\mathbf{B}$, we fix all the SVM weights $\{\mathbf{w}_j^x\}_{j=1}^m$ and $\{\mathbf{w}_k^b\}_{k\in\mathcal{Y}}$ in (1). This results in an integer programming problem and is therefore NP-hard. As a result, we take an approach akin to the discrete cyclic coordinate descent algorithm in [18]. We learn binary codes column by column in $\mathbf{B}$ iteratively based on the rest of the binary codes until convergence. For updating bit j in $\mathbf{b}_i$ for $i = 1,\dots,n$, we reduce (1) to the following optimization problem:

$$\begin{aligned} \min_{\{b_{ij}\}_{i=1}^n} \quad & \sum_{i=1}^n L(b_{ij},i,j) + \gamma \left| \sum_{i=1}^n b_{ij} \right| \\ \text{s.t.} \quad & \forall i \quad b_{ij} \in \{-1,+1\} \end{aligned} \tag{2}$$

where $L(z,i,j)$ represents the total hinge losses depending specifically on data point i when the corresponding bit j is equal to $z \in \{-1,+1\}$, specifically:

$$L(z,i,j) = \beta C^b \max_{k\in\mathcal{Y}} \big(\mathbb{1}[y_i \neq k] + z(w^b_{k,j} - w^b_{y_i,j}) + \theta_{ijk}\big) + C^x \max\big(0, 1 - z(\mathbf{w}_j^x)^\top \varphi(\mathbf{x}_i)\big)$$

where $\theta_{ijk} = \sum_{u\neq j} b_{iu}(w^b_{k,u} - w^b_{y_i,u})$ is a bias term that depends on the binary codes excluding bit j. In case of no penalty for imbalance ($\gamma = 0$), the closed-form solution is $b_{ij} = \arg\min_{z\in\{-1,+1\}} L(z,i,j)$ for $i = 1,\dots,n$. In our case, we need to consider the trade-off between the hinge losses and the imbalance penalty. This is achieved by sorting the data points with respect to the difference $\delta_{ij} = L(-1,i,j) - L(+1,i,j)$ and finding a cutting location between $+1/-1$ assignments as follows. Let I contain the indices of data points in ascending order, then the cutting location is determined by

$$\text{cut} = \arg\min_{l\in\{0,\dots,n\}} \left(\gamma|2l-n| + \sum_{i=1}^{l} L\big(-1, I[i], j\big) + \sum_{i=l+1}^{n} L\big(+1, I[i], j\big) \right). \tag{3}$$

Next, we assign -1 to bit j of data points from $I[1]$ to $I[\text{cut}]$ and $+1$ for data points from $I[\text{cut}+1]$ to $I[n]$. Note that the first sum in (3) represents the cumulative hinge losses for the

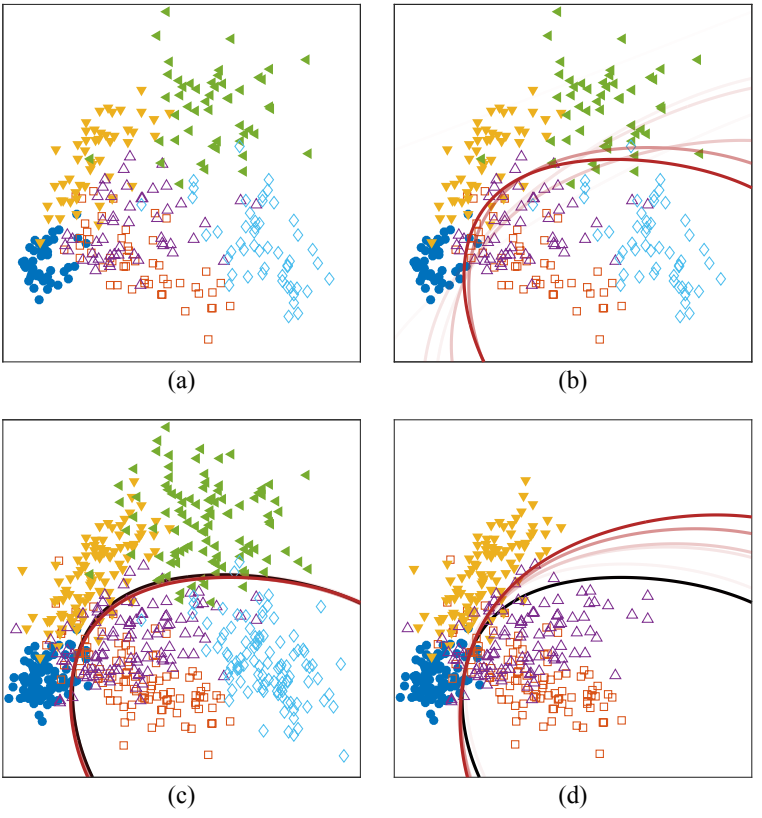

Figure 1: (a) A sample dataset of 300 data points from 6 classes represented by colors and shapes. Assignments of $+1$ and -1 are indicated by filled and empty shapes, respectively. (b) Each red line represents a hyperplane corresponding to the weight vector at each step of the OCAS method *i.e.* $(\mathbf{w}_t^b)^\top \mathbf{x} = 0$. Increasing transparency of the lines indicates earlier iterations of the execution *i.e.* smaller t. (c) The number of data points is increased to 600 by adding new points from the same distributions in (b). Red lines represent the hyperplanes when the OCAS method is initialized with the solution at (b) shown by a black line. (d) Two classes are deleted from (c). Red lines represent the hyperplanes when the OCAS method is initialized with the solution at (c) shown by a black line.

assignments of -1 and the second one represents the same for $+1$ assignments. We repeat these column updates for $j = 1,\ldots,m$ until $\mathbf{B}$ converges; that typically requires a few full updates of $\mathbf{B}$. The proposed Supervised Incremental Hashing (SIH) method is summarized in Algorithm 1.

2.6 Incremental Updates to Hash Functions

Our method can be efficiently updated whenever a modification occurs in the training set. We describe the initialization of our incremental learning strategy from the easiest to the hardest type of modification as follows:

1. *Deleting existing classes*: Let $\mathcal{Y}_d$ be the set of deleted class labels, we remove $\mathbf{w}_k^y$

Algorithm 1 Supervised Incremental Hashing

Input: Training data $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, code length m, number of anchor points r, maximum iteration number max_iter, parameters C^x, C^b, λ and γ.

Output: Binary codes $\mathbf{B}$, hash function $H(\mathbf{x}) = \text{sign}\big((\mathbf{W}^x)^\top \mathbf{x}\big)$.

Randomly select r anchor points $\{\mathbf{a}_i\}_{i=1}^s$ from the training data
Compute the kernel function $\varphi(\mathbf{x}_i)$ for $i = 1, \ldots, n$.
Initialize binary codes $\mathbf{B}^{(0)}$ with a random string in $\{+1, -1\}^m$ for each class.
for $t \leftarrow 1$ to max_iter **do**
  for $j \leftarrow 1$ to m **in parallel do**
    if column j of $\mathbf{B}^{(t-1)} \neq$ column j of $\mathbf{B}^{(t)}$, **then**
      Train binary SVM for bit j using the OCAS algorithm on $\big\{\big(\varphi(\mathbf{x}_i), b_{ij}^{(t-1)}\big)\big\}_{i=1}^n$.
    end if
  end for
  Train linear multiclass SVM using the OCAM algorithm where the data is $\big\{\big(\mathbf{b}_i^{(t-1)}, y_i\big)\big\}_{i=1}^n$.
  repeat
    for j from 1 to m **do**
      Compute cut in (3) of the main paper
      Update column j in $\mathbf{B}^{(t)}$ according to cut and I.
    end for
  until $\mathbf{B}^{(t)}$ convergences
  if $\mathbf{B}^{(t-1)} = \mathbf{B}^{(t)}$ **then**
    break the loop
  end if
end for

for $k \in \mathcal{Y}_d$ from the multi-class SVM $\{\mathbf{w}_k^b\}_{k\in\mathcal{Y}}$ and the rows of $\mathbf{B}$ that correspond to images associated with any class in $\mathcal{Y}_d$ and start our incremental strategy to adopt to the changes.

2. *Adding images to existing classes*: For each class, the most frequent binary string pattern from $\mathbf{B}$ is found. Next, each new image is initialized with the corresponding binary code according to its class and the hashing functions are changed incrementally with our method.

3. *Adding new classes*: New rows to $\mathbf{B}$ are added corresponding to new images, and we initialize each new class with a random binary string. Since we do not have multi-class SVM trained on the new classes, we only train the multi-class SVM $\{\mathbf{w}_k^b\}_{k\in\mathcal{Y}_f}$ on $\mathbf{B}_f$ from scratch where the subscript f denotes the final version. The hashing functions are learned with our incremental algorithm.

3 Experiments

We conducted extensive experiments to assess the effectiveness and efficiency of the proposed method, SIH. Before evaluating the performance on dynamic datasets, we first compared our method with the state-of-the-art supervised hashing techniques including CCA-ITQ [8],

KSH [15], FastHash [13], and SDH [18]. The performance of the methods was analyzed regarding search accuracy and training/testing time on three large-scale datasets – CIFAR-10, MNIST, and NUS-WIDE. All experiments were performed in the MATLAB environment on a machine with a 2.8 GHz Intel Core i7 CPU and 16GB RAM using the public code provided by the authors with their suggested parameters unless otherwise specified.

3.1 Datasets and Experimental Setup

The CIFAR-10 dataset [9] has 60,000 labeled images from 10 classes of vehicles and animals. There exist 50,000 training and 10,000 test images in the dataset. In our experiments, each image is represented by a GIST descriptor [16] of 512 dimensions. The MNIST dataset [12] is a collection of 28×28 pixel images of handwritten digits. The dataset contains a training set of 60,000 examples and a test set of 10,000 examples. The NUS-WIDE dataset [2] includes 269,648 images and associated semantic labels of 81 concepts from Flickr. Unlike the other two datasets, each image is associated with zero or more labels. Therefore, we only used the training images that are associated with exactly one of the most frequent 10 tags. The resultant training set has 38,255 images with 10 classes. Similarly, we constructed a query set of 25,386 images among the test images associated with only one of those tags. A 500-dimensional bag-of-words vector is used to represent each image in our experiments where the codebook is generated from SIFT descriptors.

For the KSH method, 5,000 images are sampled uniformly from the training part of each dataset because of long training time of this technique (Table 1). The tree depth parameter of FastHash is set to 2 due to its higher computational complexity. We sample 1,000 images uniformly from the training sets as anchor points $\{\mathbf{a}_l\}_{l=1}^{s}$ for all the hashing methods with kernels (SIH, KSH, and SDH). All these kernel methods used an RBF kernel $\varphi(\mathbf{x},\mathbf{a}) = \exp\left(-\|\mathbf{x}-\mathbf{a}\|^2/2\sigma^2\right)$ with a kernel width σ which is adjusted for each dataset using cross-validation. For all datasets, features are first centered to zero; then each data sample is normalized to have the unit length.

For SIH, we adjusted the soft margin parameters of the SVMs and the kernel width based on cross-validation on some binary codes obtained from other hashing methods. Next, we empirically adjusted the other parameters λ and γ for each dataset. We prefer the multi-class SVM dominate the process of learning binary codes by setting λ to a large number. For example, for the CIFAR-10 dataset the parameter values are $C^x = 16$, $C^b = 10^{-3}$, $\lambda = m\times10^8$ and $\gamma = 10^5$. The maximum number of iterations is set to 5 like SDH.

3.2 Evaluation Methodology

We employed leave-one-out validation on test data in our experiments to assess the quality of hash functions learned from training data. We used each image in the test data as a query while we treat the remaining test images as a retrieval set. For each query, we rank images according to Hamming distances between their binary codes and that of the given query. For quantitative analysis, we compute precision and recall value at each Hamming radius. Next, these values are used for computing mean average precision (mAP) and also mean precision values.

3.3 Effects of Training and Anchor Set Size

We start our experiments by analyzing the effects of training and anchor set size on retrieval. The SIH method is analyzed in details with different numbers of anchor points and training points for comparison (Table 1). As expected, larger training and anchor sets provide better performance with longer execution time. The number of anchor points has a greater influence on retrieval performance than training size. Note that the test time is affected by the scale of the anchor set. Our method outperforms other methods in retrieval performance while it has competitive execution time. Note that the training time in Table 1 can be enhanced by a computer with a larger number of nodes as our method has a distributed inference algorithm.

Table 1: Our method (SIH) is compared in terms of mean average precision (mAP), mean of precision at Hamming radius $r = 2$, training and test time for 32-bit hash codes on the CIFAR-10 dataset. The experiments were performed on a machine with an Intel quad-core processor.

Method	Training Set Size	Anchor Set Size	mAP	Precision at $r = 2$	Training Time (s)	Test Time (μs)
SIH	5,000	300	0.322	0.383	2.7	6.1
	5,000	1,000	0.366	0.420	7.5	15.9
	5,000	3,000	0.393	0.438	47.3	45.7
	50,000	300	0.360	0.417	61.7	6.3
	50,000	1,000	**0.429**	0.478	171.1	17.7
SDH	50,000	1,000	0.422	**0.478**	25.0	24.9
Fasthash	50,000	-	0.357	0.271	628.4	102.6
KSH	5,000	1,000	0.343	0.322	2,802.8	26.8
CCA-ITQ	50,000	-	0.325	0.410	1.5	3.8

3.4 Retrieval Performance Analysis

We report retrieval performance as mean average precision (mAP) for all methods on the three datasets in Figure 2 for different code lengths. Precision-recall curves of all methods on all datasets are displayed in Figure 3 for 32-bit length codes. As seen in the table and figures, the proposed hashing method outperforms the state-of-the-art supervised hashing methods on the CIFAR-10 and MNIST datasets, in which each image belongs to a single class. In addition, the imbalance penalty improves the retrieval performance for almost all cases. Note that we used a sequential sampling procedure for initializing $\mathbf{B}$ where balanced binary codes have a higher probability for the SIH version with imbalance penalty *i.e.* $\gamma > 0$. SDH provides the second best performance. On the other hand, our method has the best performance on the NUS-WIDE dataset for shorter binary codes, and it has competitive results along with FastHash and SDH for longer codes.

3.5 Retrieval Performance Analysis for Dynamic Datasets

We evaluated our incremental hashing strategy in comparison to the passive and from-scratch strategies for the aforementioned three types of modifications on the CIFAR-10 dataset with

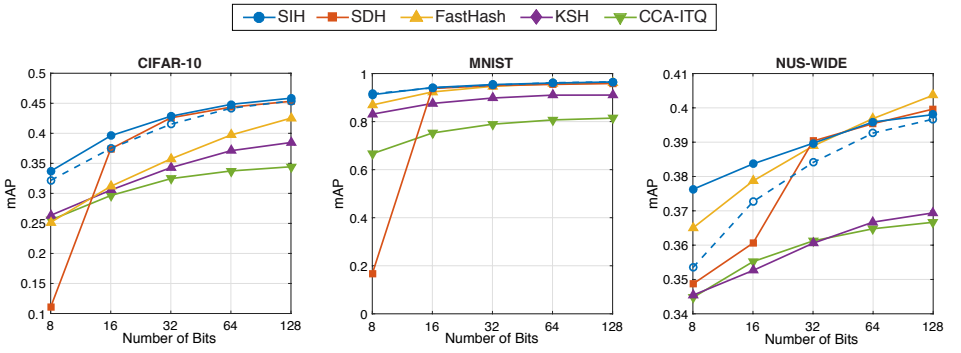

Figure 2: Our method (SIH) is compared with the state-of-the-art methods on, from left to right, the CIFAR-10, MNIST and NUS-WIDE datasets in terms of mean average precision (mAP). Dashed line represents SIH without imbalance penalty.

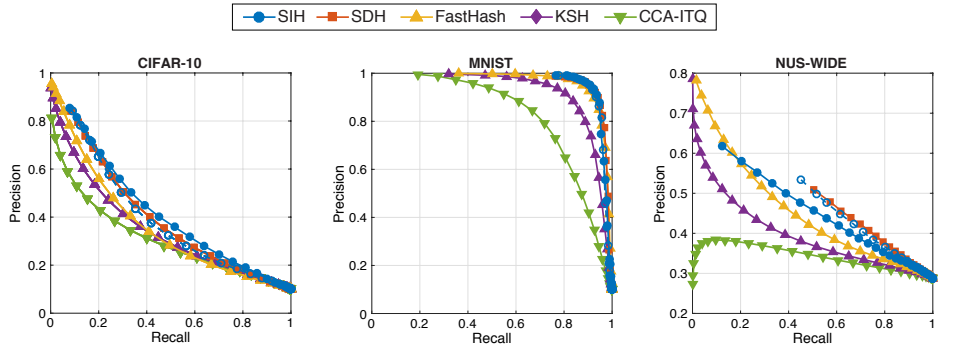

Figure 3: Our method (SIH) is compared with the state-of-the-art methods on, from left to right, the CIFAR-10, MNIST and NUS-WIDE datasets by precision-recall curves for 32-bit length hash codes. Dashed line represents SIH without imbalance penalty.

32 bits binary codes in terms of mAP scores and training time. For the case of deleting classes, we first learned hash functions from the entire training data. Next, we removed images from randomly selected classes and learned hash functions both from scratch and incrementally on the final training data. We repeated this 25 times and reported mAP and training time. A similar methodology in the reverse order was employed for adding new classes case. For adding new images to existing classes case, we first learned hash functions from a randomly selected subset of training data and then used the entire training set for from-scratch and incremental strategies. For adding new classes and new images to existing classes, we used the same query set as in the previous section. For deleting classes case, we removed the images associated with those classes and computed mAP for the rest of the queries. The mAP and training time are shown in Figures 4-6 for three types of modifications on all datasets. Note that the training time reported for the passive strategy represents the initial computation time, and it does not have any additional computation for changes in the dataset. Our incremental hashing strategy reaches the same retrieval performance as the from-scratch strategy while requiring shorter training time. Note that the total computation time for both training on 10% of training data and incrementally training on the entire data is less than the training time on

the entire data set from scratch. As a result, it might be a sound strategy to use our supervised hashing method as a combination of training on a small set of representative images followed by incremental training on the entire training data. This will reduce the training time of our method in Table 1.

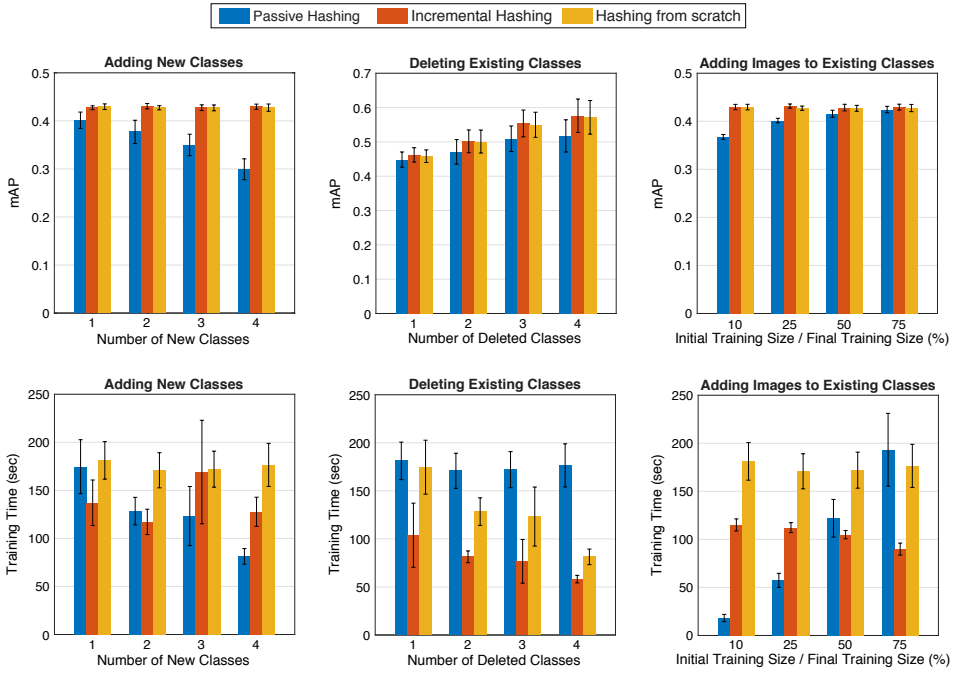

Figure 4: Incremental Hashing is compared with from-scratch and passive hashing for different types of modifications, from left to right, adding new classes, deleting existing classes and adding new images to existing classes on CIFAR-10 in terms of mean average precision (mAP) in the first row and in terms of training time in the second row at 32-bits.

4 Conclusion

We presented an incremental supervised hashing method based on a two-stage classification framework. We formulated a joint optimization task for the classification problem and the problem of finding optimal binary codes. An efficient algorithm was developed for learning hash functions in a distributed scheme where the sub-problems are solved independently. Experiments validate that the incremental hashing strategy for dynamic datasets is capable of updating hash functions efficiently. Besides, the proposed approach provides higher quality codes with well-balanced bits and better generalization.

References

[1] R. Batuwita and V. Palade. *Class Imbalance Learning Methods for Support Vector Machines*, pages 83–99. John Wiley & Sons, Inc., 2013.

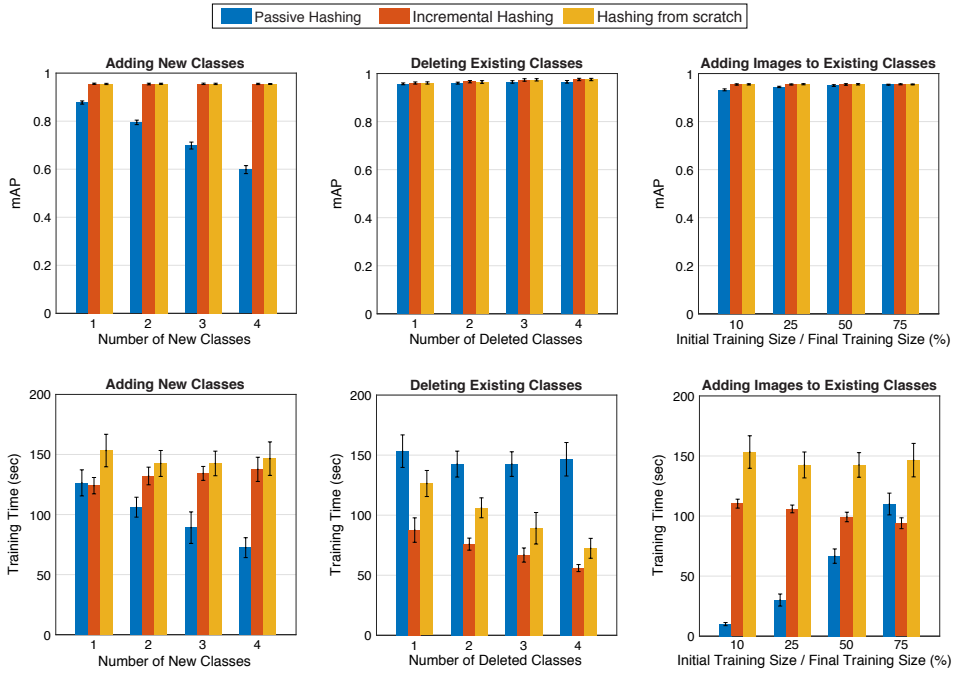

Figure 5: Incremental Hashing is compared with from-scratch and passive hashing for different types of modifications, from left to right, adding new classes, deleting existing classes and adding new images to existing classes on MNIST in terms of mean average precision (mAP) in the first row and in terms of training time in the second row at 32-bits.

[2] T.-S. Chua, J. Tang, R. Hong, H. Li, Z. Luo, and Y.-T. Zheng. Nus-wide: A real-world web image database from national university of singapore. In *CIVR*, Santorini, Greece, 2009.

[3] K. Crammer and Y. Singer. On the algorithmic implementation of multiclass kernel-based vector machines. *JMLR*, 2:265–292, March 2002.

[4] V. Franc and S. Sonnenburg. Optimized cutting plane algorithm for large-scale risk minimization. *JMLR*, 10:2157–2192, December 2009.

[5] A. Gionis, P. Indyk, and R. Motwani. Similarity search in high dimensions via hashing. In *VLDB*, pages 518–529, 1999.

[6] Y. Gong, S. Lazebnik, A. Gordo, and F. Perronnin. Iterative quantization: a Procrustean approach to learning binary codes for large-scale image retrieval. *TPAMI*, 35(12):2916–2929, December 2013.

[7] J. He, R. Radhakrishnan, S.-F. Chang, and C. Bauer. Compact hashing with joint optimization of search accuracy and time. In *CVPR*, pages 753–760, 2011.

[8] R. He, Y. Cai, T. Tan, and L. S. Davis. Learning predictable binary codes for face indexing. *Pattern Recognition*, 48(10):3160 – 3168, 2015.

[9] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

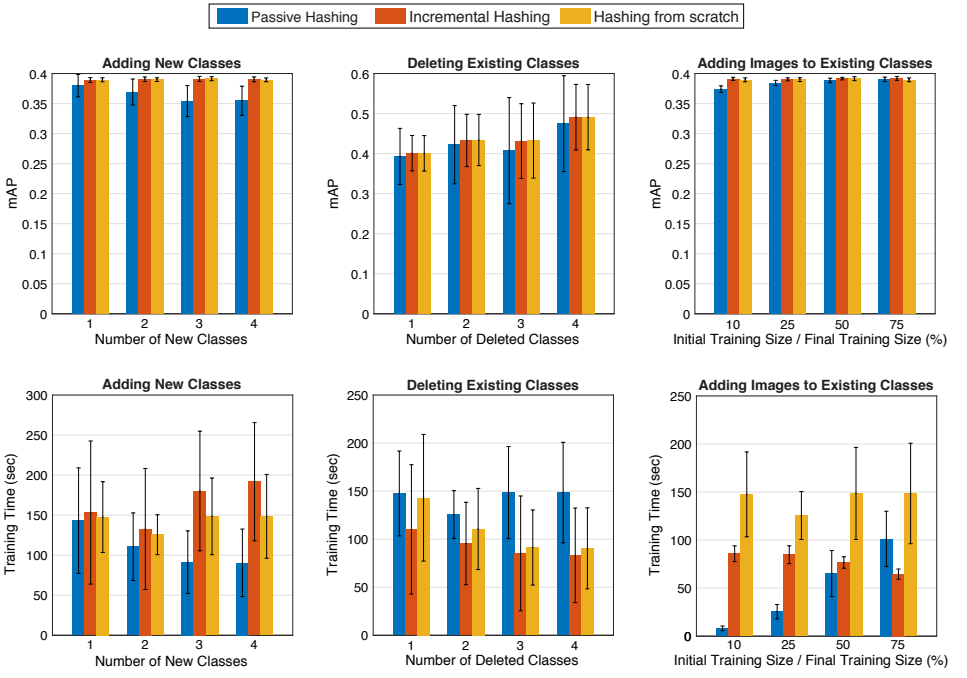

Figure 6: Incremental Hashing is compared with from-scratch and passive hashing for different types of modifications, from left to right, adding new classes, deleting existing classes and adding new images to existing classes on NUS-WIDE in terms of mean average precision (mAP) in the first row and in terms of training time in the second row at 32-bits.

[10] B. Kulis and T. Darrell. Learning to hash with binary reconstructive embeddings. In *NIPS*, pages 1042–1050, 2009.

[11] B. Kulis and K. Grauman. Kernelized Locality-Sensitive Hashing. *TPAMI*, 34(6):1092–1104, 2011.

[12] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, Nov 1998.

[13] G. Lin, C. Shen, Q. Shi, Anton V. d. Hengel, and D. Suter. Fast supervised hashing with decision trees for high-dimensional data. In *CVPR*, pages 1971–1978, June 2014.

[14] W. Liu, J. Wang, S. Kumar, and S.-F. Chang. Hashing with graphs. In *ICML*, pages 1–8, 2011.

[15] W. Liu, J. Wang, R. Ji, Y.-G. Jiang, and S.-F. Chang. Supervised hashing with kernels. In *CVPR*, pages 2074–2081, 2012.

[16] A. Oliva and A. Torralba. Modeling the shape of the scene: A holistic representation of the spatial envelope. *IJCV*, 42(3):145–175, 2001.

[17] M. Rastegari, A. Farhadi, and D. Forsyth. Attribute discovery via predictable discriminative binary codes. In *ECCV*, pages 876–889, 2012.

[18] F. Shen, C. Shen, W. Liu, and H. Tao Shen. Supervised discrete hashing. In *CVPR*, June 2015.

[19] C.-H. Tsai, C.-Y. Lin, and C.-J. Lin. Incremental and decremental training for linear classification. In *International ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 343–352, New York, NY, USA, 2014.

[20] Yair Weiss, Antonio Torralba, and Rob Fergus. Spectral hashing. In *NIPS*, pages 1753–1760, 2009.

Appendix

The SVM problem can be described as an unconstrained regularized risk minimization problem as follows:

$$\mathbf{w}^* = \arg\min_{\mathbf{w}\in\mathbb{R}^d} F(\mathbf{w}) := \frac{1}{2}\|\mathbf{w}\|^2 + CR(\mathbf{w}), \tag{4}$$

where $\mathbf{w} \in \mathbb{R}^d$ denotes the weight vector to be learned, $\frac{1}{2}\|\mathbf{w}\|^2$ is a quadratic regularization term, $C > 0$ is a fixed regularization constant and $R : \mathbb{R}^d \to \mathbb{R}$ is a non-negative convex risk function *e.g.* hinge loss on training data, $R(\mathbf{w}) = \sum_{i=1}^{n} \max(0, 1 - y_i \mathbf{w}^\top \mathbf{x}_i)$.

The cutting plane method approximates the convex function $F(\mathbf{w})$ by a piecewise linear function $F_t(\mathbf{w})$ with t cutting planes. Let $\mathbf{w}_t$ and $\mathbf{w}_t^b$ be the *current solution* and the *best-so-far solution* of the OCAS method at iteration t, respectively. The stopping condition of the OCAS method is defined for a given tolerance parameter ε as $F(\mathbf{w}_t^b) - F_t(\mathbf{w}_t) \leq \varepsilon$. Given in [4], assume that $\|\partial R(\mathbf{w})\| \leq G$ for all $\mathbf{w} \in \mathbb{R}^d$, and $F(\mathbf{w}_t^b) - F_t(\mathbf{w}_t) = \varepsilon_t > 0$, then

$$\varepsilon_t - \varepsilon_{t+1} \geq \frac{\varepsilon_t}{2} \min\left(1, \frac{\varepsilon_t}{4C^2G^2}\right). \tag{5}$$

From (5), for any $\varepsilon_t > 4C^2G^2$, we say that $\varepsilon_{t+1} \leq \frac{\varepsilon_t}{2}$; starting from a weight vector $\mathbf{w}_0$, we can reach at a level precision better than $4C^2G^2$ after at most $\log_2 \frac{F(\mathbf{w}_0)}{4C^2G^2}$ iterations [4]. Subsequently, we can find the remaining number of iterations by solving the following differential equation:

$$\varepsilon_{t+1} + \frac{\varepsilon_t^2}{8C^2G^2} - \varepsilon_t = 0. \tag{6}$$

Franc and Sonnenburg provide a solution to this equation in [4]. We need $\frac{8C^2G^2}{\varepsilon} - 2$ more iterations until convergence. The OCAS method is initialized its computation from the zero vector. Therefore, the total number of iterations of the OCAS method to solve an SVM problem is at most

$$\log_2 \frac{F(\mathbf{0})}{4C^2G^2} + \frac{8C^2G^2}{\varepsilon} - 2. \tag{7}$$

In Supervised Incremental Hashing, we use SVMs with hinge losses. Thus, $F(\mathbf{0})$ is equal to nC for a dataset with n images. The initial weight vector $\mathbf{w}_0$ has no effect on the number of iterations following a precision level better than $4C^2G^2$. However, considering an initial vector $\mathbf{w}_0$ in our incremental setting such that $F(\mathbf{w}_0) < nC$ *i.e.* having a better solution $\mathbf{w}_0$ than the zero vector, we can reduce the number of iterations until achieving such a precision level by at most

$$\log_2 \frac{F(\mathbf{0})}{4C^2G^2} - \log_2 \frac{F(\mathbf{w}_0)}{4C^2G^2} = \log_2 \frac{nC}{F(\mathbf{w}_0)}. \tag{8}$$

This concludes the claim in Section 2.